Toward Improved Landing Precision on Mars

Aron A. Wolf
Aron.A.Wolf@jpl.nasa.gov
818-354-6917

Behcet Acikmese
818-393-7749
Behcet.Acikmese@jpl.nasa.gov

Yang Cheng
818-354-1857
Yang.Cheng@jpl.nasa.gov,

Jordi Casoliva
818-354-7590
Jordi.Casoliva@jpl.nasa.gov

John M. Carson
818-354-4340
john.m.carson@jpl.nasa.gov,

Mark C. Ivanov
818-393-2442
Mark.C.Ivanov@jpl.nasa.gov

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109

Abstract— Mars landers to date have flown ballistic entry trajectories with no trajectory control after the final maneuver before entry. [12]Improvements in landing accuracies (from ~150 km from the target for Mars Pathfinder to ~30-40 km for MER and Phoenix) have been driven by approach navigation improvements. MSL will fly the first guided-entry trajectory to Mars, further improving accuracy to ~10-12 km from the target.

For future missions, landing within ~100m is desired to assure landing safety close to a target of high scientific interest in irregular terrain, or to land near a previously landed asset. Improvements in approach navigation alone are not sufficient to achieve this requirement. If approach navigation error and IMU error are eliminated, the dominant error source is wind drift on the parachute, with map-tie error also significant. Correcting these errors requires terrain-relative navigation (TRN), which can be accomplished with passive imaging supplemented by radar for terrain sensing (with onboard navigation capable of processing measurements from IMU, imaging, and radar). Additionally, near-optimal-ΔV powered descent guidance is needed to minimize the amount of propellant required to reach the target. The capability to land within 100m can be applied in different landing modes depending on how much fuel is carried.

[1] 978-1-4244-7351-9/11/$26.00 ©2011 IEEE
[2] IEEEAC paper#1209, Version 1, Updated 2010:10:26

TABLE OF CONTENTS

1. INTRODUCTION

Error sources affecting delivery uncertainty (the dimensions of the "landing ellipse") have been discussed in detail in [3]. Dispersions at engine ignition must be corrected during powered descent. Consequently, these dispersions drive propellant consumption (and propellant penalty compared to landing modes on MSL and previous missions), which makes the magnitude of these dispersions a critical metric of interest.

Mars landers to date have all used a non-steerable parachute to reduce speed before touchdown. For any lander using this

architecture, wind drift on the parachute is a major source of delivery error. When other error sources are substantially reduced (e.g. attitude, position, and velocity knowledge are improved at entry and atmosphere and aerodynamics modeling errors are "flown out" using entry guidance and bank control) wind drift becomes the "tall tent pole", responsible for delivery errors of up to ~3km. Wind drift can be corrected during powered descent; however large propellant expenditures (~65 kg per km for an MSL-class vehicle) are required.

Pinpoint Landing is defined as landing within 100m of a single preselected target (the "center of the ellipse"), as illustrated in the schematic shown in Figure 1.

For missions where the propellant mass penalty of flying to the center of the ellipse is unacceptably high, another landing mode, which may be called "Multipoint", can be used. In Multipoint, the fuel quantity onboard the spacecraft is not sufficient to fly from every point in the landing ellipse to the center of the ellipse; however the lander does carry enough fuel for diverting to reach a safe landing site. Orbital imagery is used to preselect a number of safe sites throughout the landing ellipse, to assure that there is always at least one safe spot within fuel range at engine ignition. The number of preselected targets required depends on the amount of propellant carried and the size of the ellipse at ignition. The lander needs decisionmaking logic onboard to select a target from the preselected target set after beginning descent imagery. The "cartoon" in Figure 1 shows a comparison of Pinpoint Landing and Multipoint.

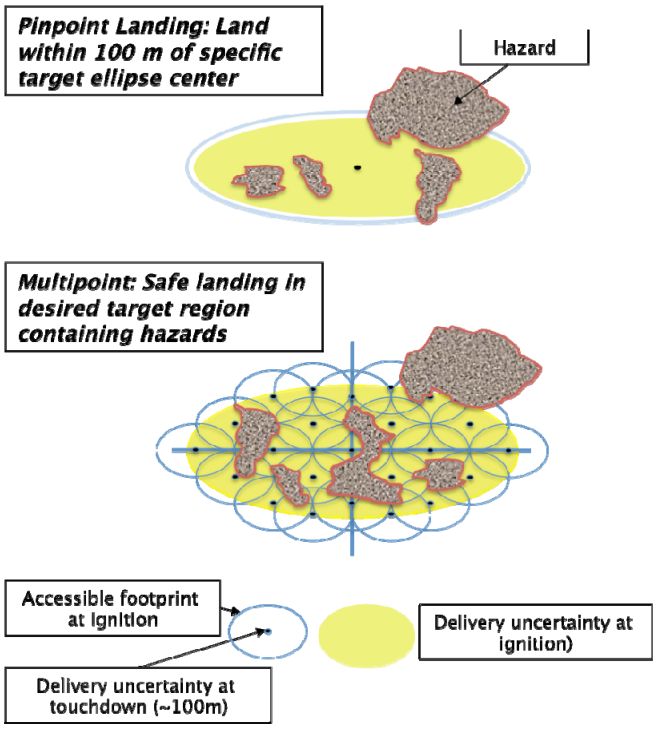

Figure 1. Capability to deliver within 100m of a target can be applied in both Pinpoint Landing (upper) and Multipoint Landing (lower)

In both Pinpoint and Multipoint Landing scenarios, the technologies described above (TRN, near-optimal-ΔV powered descent guidance, and onboard navigation processing all measurement types) are used to locate the spacecraft relative to terrain features and fly it to the desired target. There are two differences between the two scenarios: 1) the lander carries less fuel in Multipoint, and 2) in Multipoint, the lander needs additional decisionmaking logic onboard to selects its target from the pre-selected set of targets after taking its initial descent imagery.

An alternative landing mode using these technologies could be labeled "Multi-hazard". In this approach, which is the "mirror image" of the Multipoint approach, the lander uses its TRN capability to avoid areas pre-designated as hazardous. If the lander finds itself on its way to a landing within the contours of a hazardous area, a minimum-fuel divert is flown to the nearest point outside the hazardous area contour.

2. PROPELLANT PENALTY

A useful metric for the propellant requirement is Propellant Mass Fraction (PMF), defined as propellant mass / total wet mass at ignition (where total wet mass at ignition = propellant mass at ignition + dry mass). PMF is a useful parameter because the PMF required is independent of the mass of the lander. Consequently, the same mass fraction applies as dry mass is scaled up or down.

For an MSL-class spacecraft, propellant load is ~400kg and wet mass at ignition ~2100 kg, yielding PMF = ~19%. Each percent PMF is equivalent to ~21 kg of propellant. Simulations with the ΔV-optimal powered descent guidance algorithm described below have shown that each additional km of maneuvering capability requires an increase in propellant equivalent to ~3% PMF (=~63 kg for an MSL-class spacecraft).

Ways to mitigate or reduce dispersions at ignition include use of a range trigger for parachute deployment (for which there is an associated altitude penalty) and potential improvements in atmospheric entry guidance. Use of a steerable parachute has a;lso been considered; however, in order to be usable on a Mars lander mission, a steerable parachute would have to be designed to divert without turning, due to the excessive time and altitude loss required to accomplish a turn in the thin Martian atmosphere. This could be accomplished by opening vents or similar devices to produce sideslip. Control authority and stability remain unresolved issues.

3. POWERED DESCENT GUIDANCE

The powered descent phase of Mars EDL is a short-duration event requiring guidance algorithms that provide valid solutions in a minimal amount of time. These Powered-Descent Guidance (PDG) algorithms must minimize landing error while simultaneously satisfying the governing physics and the physical state and control constraints that affect the landing vehicle performance. Traditional PDG methods have implemented point designs that provide limited divert capability and do not account for target direction; with these algorithms, propellant mass penalties are unacceptably high for the large diverts of several km required for Pinpoint or Multipoint landing [8]. Further, these methods do not guarantee satisfaction of the state or control constraints outside of the point design. The development of an exhaustive database of handcrafted point designs in place of onboard trajectory computation is infeasible, as it creates a Validation and Verification (V&V) challenge and increases the cost of ensuring satisfaction of state and control constraints for each handcrafted guidance design.

To truly enable Pinpoint and Multipoint Landing, a new PDG algorithm must be developed that by design generates guidance profiles that guarantee satisfaction of the state and control constraints. If this algorithm is to be run onboard, it must be numerically efficient and must guarantee finding solutions given limited time and computational resources. The state constraints in PDG include limitations on the position and velocity attained from application of the guidance law. Position constraints include avoiding subsurface flight or ensuring descent within a prescribed glide-slope cone, and velocity constraints include the maximum speed to avoid excessive drag or the generation of aerodynamic shock waves (assuming the lander is not designed to handle them). Additionally, control constraints must be applied to ensure that physical thrust limitations are not exceeded and that thrust-pointing constraints are not violated. After powered descent ignition, typical liquid thrusters used for landing cannot be throttled fully off, so guidance algorithms must account for both a maximum and a minimum thrust bound in the computation of valid thrust profiles. Further, onboard sensors for TRN typically have field-of-view limitations that put an overall pointing requirement on the landing vehicle attitude. This requirement places pointing limits on the allowable thrust-vector directions for translational thrust.

The aforementioned state and control constraints have been incorporated into a numerically efficient guidance algorithm which minimizes landing error by posing the problem as a type of convex optimization problem known as a Second order Cone Program (SoCP) [12]. The SoCP formulation enables two alternative onboard solution methods. The first is to use numerically efficient interior-point solvers that have deterministic stopping criteria, and which guarantee that if a feasible solution exists, the solver will find the globally optimal solution [11] in a predetermined number of mathematical operations. This guarantee provides robustness, ensuring that the algorithm converges on the solution when the spacecraft is capable of performing a powered-descent maneuver under the specified constraints. The second approach is to use an onboard table of pre-computed optimal trajectories, along with an in-flight convex interpolation procedure, to generate nearly-optimal solutions for most of the feasible space. The advantage of the table lookup approach is that it avoids using onboard optimization, and the onboard convex interpolation procedure reduces the size of the table while still ensuring constraint satisfaction. This second approach also provides a type of onboard database but avoids the time-consuming generation and V&V of many times more handcrafted trajectories. We anticipate that a mission would choose between the two approaches depending on its specific requirements.

The convex PDG algorithm was originally developed in [8] for fuel-optimal Pinpoint Landing guidance with explicitly enforced state constraints and minimum and maximum thrust bounds. The numerical efficiency of this convex formulation is considered a strong candidate for future onboard flight implementation [14]. The convex Pinpoint Landing PDG work was further extended in [12] for Multipoint Landing to handle situations where fuel limitations prevent the ability to fly all the way to the designated target. In such cases, the algorithm instead minimizes the landing error while ensuring satisfaction of the state and control constraints.Thrust pointing constraints have also been included in the convex PDG algorithm framework [13].

For a representative spacecraft performing a powered descent to a landing on Mars, the simulations in [13] highlight the capabilities of the convex PDG algorithm to generate thrust profiles that enable large diverts (See Figure 2) to minimize landing dispersions while simultaneously enforcing the state constraints on position and velocity and the control constraints on the thrust (throttle and pointing angle).

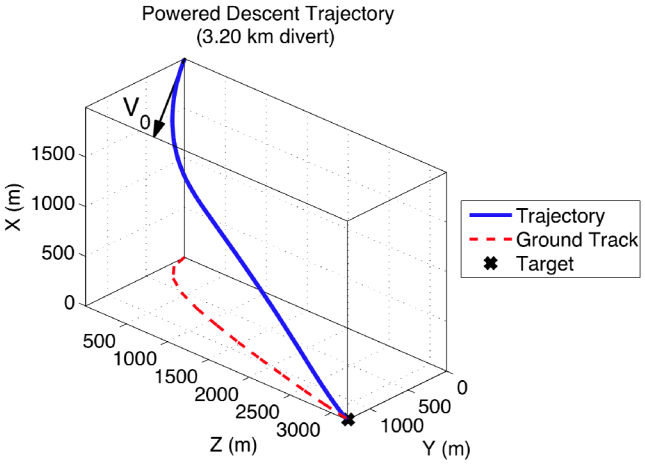

Figure 2. Convex PDG algorithm provides guidance

profiles that enable large diverts while satisfying state and control constraints [13].

4. HAZARD DETECTION USING ORBITAL IMAGERY

With imagery taken by the HIRISE camera onboard the Mars Reconnaissance Orbiter (MRO), it is possible to map rocks down to 1.5m in diameter. For the purposes of hazard mapping, rocks on Mars can be modeled as "hemispheres", thus the height of a 1.5m-diameter rock is ~0.75m.

However, it is not necessary to detect every individual rock of height greater than the lander's hazard tolerance in order to be able to achieve an acceptably high probability of safe landing. Sites on Mars and Earth show rock size-frequency distributions that follow an exponential when expressed in cumulative fractional area covered by rocks of a given diameter or larger versus diameter plots [14]. These rock abundance curves can be used to predict the probability of hitting a rock too small to be seen in orbital imagery, and are being used for this purpose on MSL. The MSL project requires the probability of landing on a rock 0.6m high / 1.2m in diameter (slightly below the threshold of detectability in HIRISE imagery) to be less than 0.5%.

Rock abundances in different terrain types follow different exponential curves. Comparison of predicted rock abundances with observed data at the MER landing sites showed close correspondence [15]. Orbital imagery shows it should be fairly straightforward to find nearly rockless plains of several km radius with slopes of < 15 deg. [17].

Nevertheless, it is possible that a site of high science value in an area with high rock abundances could be an attractive target for a future mission; or, that a future lander could be significantly less hazard-tolerant than MSL (as is the case for some proposed 2018 and MSR lander designs). The capability to detect hazards onboard in realtime could be needed to enable a lander with low hazard tolerance to detect and avoid increased abundances of rocks that are both undetectable from orbit and large enough to be hazards.

5. TERRAIN-RELATIVE NAVIGATION (TRN)

Both Pinpoint and Multipoint Landing rely on TRN technology based on computer vision. This technology recognizes the local terrain and locates the spacecraft within the local terrain frame. Currently, both active sensing (radar or lidar) and passive sensing technologies are actively pursued. Active sensing for TRN is required for night landing or for landing on bodies not as extensively mapped from orbit as Mars. The less complex passive imaging approach is well suited to daylight landings on Mars, which has been extensively mapped from orbit. In this paper we focus on passive optical TRN, which compares a descent image with an on-board reference map to locate the spacecraft during descent. Two algorithms for use in passive TRN have been previously developed at JPL: Map and Image Alignment (MAIA) and Onboard Image Reconstruction for Optical Navigation (OBIRON).

MAIA is a correlation-based TRN algorithm. In this approach, the base map is a piece of an ortho-rectified orbital image. A correlation method is used to recognize the descent image within a base map. A high pass filter is applied to both the descent image and the map to increase robustness. To accomplish the correlation, the descent image must be rectified to the orientation and scale of the base map. Then features (corners) are matched between the image and map.

OBIRON is a hybrid pattern matching and correlation approach developed originally for small body navigation. It uses reconnaissance imagery to build 3D models of surface patches. These patches become the landmarks used for navigation. During landing, the 3D patches are rendered using the current solar illumination to create landmark images (i.e. predicted images of terrain in view during descent). These landmark images are then correlated with images taken during decent to estimate offsets from the current estimate of lander position and attitude. OBIRON has been tested off-line with imagery from the NEAR and MUSES-C asteroid missions.

OBIRON and MAIA were also recently tested using imagery from ALHAT Field Test 3 that involved taking a large number of images and ancillary data during ten flights between June 20 and July 07, 2009 over Death Valley and the Nevada Test Site [18]. The flights covered a wide range of terrain types (from mountainous areas to flat lake beds), lighting conditions (from dawn to sunset) and altitude (from 1 - 8 km above ground). During each flight, six cameras (four fixed cameras and two gimbaled cameras) were used.

The MAIA algorithm was run on a total of 6,423 images, and OBIRON on a total of 3,144 images. MAIA performed well on mid-day imagery, producing position estimates for more than 97% of the images. The mean error on these position estimates was within 15 meters. However, MAIA performed poorly on imagery obtained when the sun elevation was low, because MAIA does not have a way to account for shadows cast by terrain under low-sun conditions. The algorithm produced position estimates for less than 43% of the images, and the mean error on these estimates was greater than 100 meters.

OBIRON also performed well on mid-day imagery with a success rate of 95%. However, its mean position error was about 48 meters, worse than MAIA. However, OBIRON outperformed MAIA on images taken under low-sun conditions.

This testing revealed strengths and weaknesses of both MAIA and OBIRON, pointing the way to merging MAIA and OBIRON into a unified TRN system that preserves strengths of both. We pursued the development of such an algorithm with a modular approach in order to maximize flexibility, as illustrated in Figure 3.

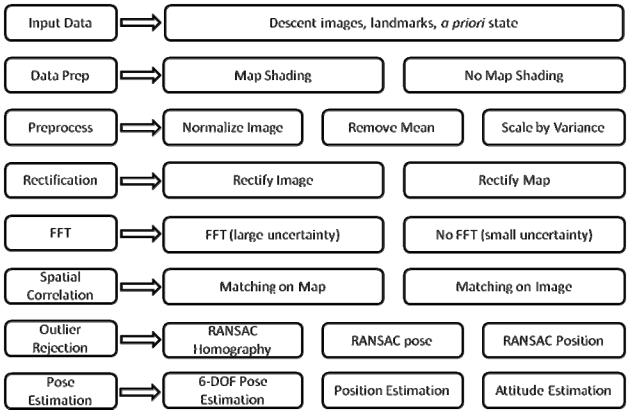

Figure 3: Modularized TRN system

The main advantage of the modularized approach is that we can assemble an optimal TRN solution for a given scenario. The two example cases below illustrate how this flexibility can be used.

Case 1: EDL takes place in early morning with sun elevation of 11°. The local terrain has moderate terrain relief. The first image is taken at 8 km altitude, at which time the assumed a priori position uncertainty of the lander is ~a few km (Figure 4). Figure 5 shows the TRN modules chosen for this situation: the map-shading module is used to handle shadow caused by the low sun angle, and the rectifying-image module is used because of high altitude and moderate terrain relief. The FFT module is used in cases like this one where a priori position uncertainty is large. Figure 6 shows the intermediate and final result.

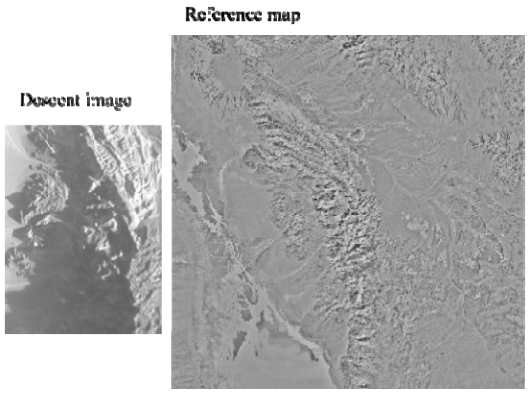

Figure 4: Descent image and corresponding reference map for example case 1

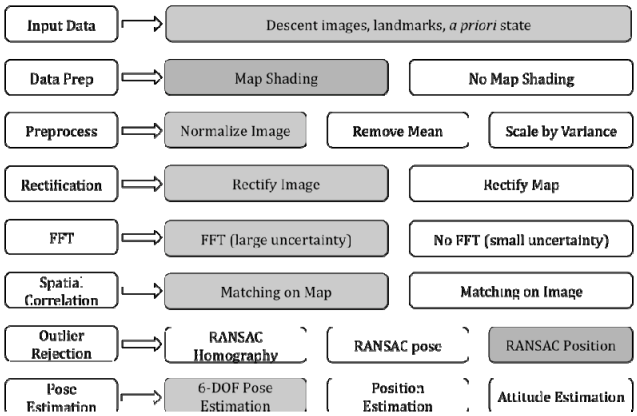

Figure 5: TRN modules used for example case 1 (shaded)

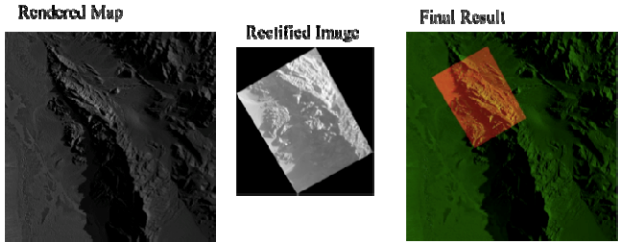

Figure 6: Shaded reference map (left), rectified descent image (center), and descent image overlaid on the reference map (right) for example case 1

Case 2: The landing occurs in mid-day in a mountainous region and the a priori spacecraft position uncertainty is very small. Figure 7 shows the TRN modules used in this case. The mid-day landing makes the map shading unnecessary, and the FFT is also not needed because the position uncertainty is small. Rectifying the map to the image frame becomes very useful because this can remove some of the terrain relief displacement. RANSAC homography outlier removal is the proper solution here. There were 156 landmarks matched, yielding a final position error of (0.06, -3.10,-4.00) in x, y, and z directions respectively.

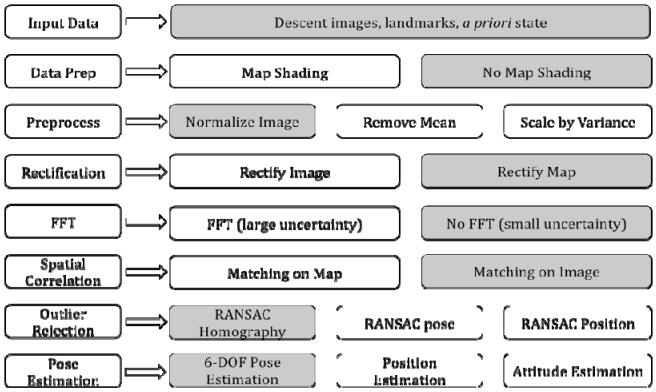

Figure 7: TRN modules used for example case 1 (shaded)

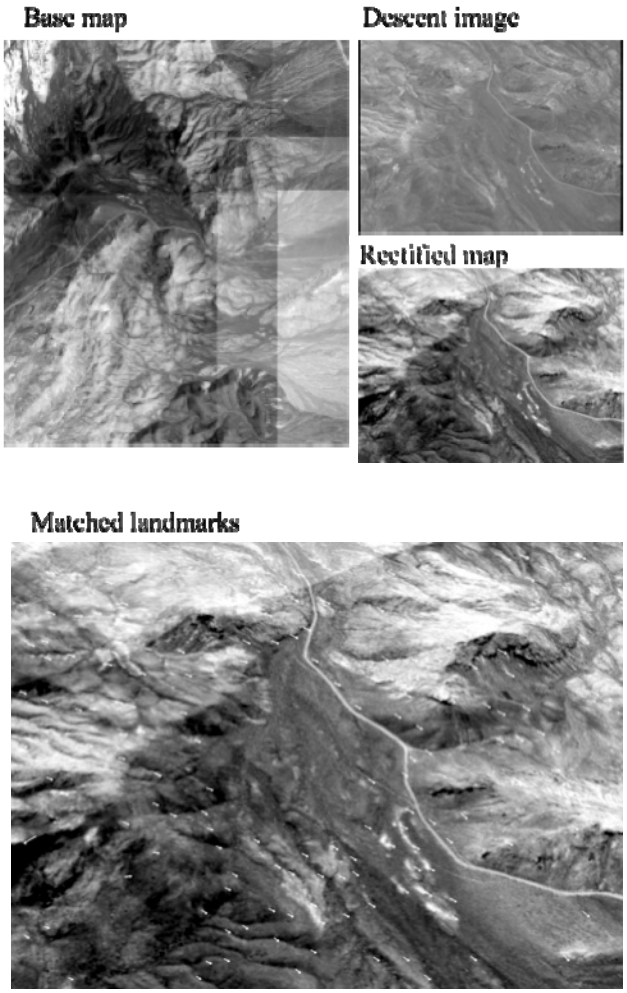

Figure 8: intermediate and final results of Case 2.

One important drawback of the TRN system developed here is that it is still too computationally intensive, as a result of which runtime in a flightlike computer is too slow to meet the probable requirements of a future mission. Future work will focus on speeding up the measurement rate via hardware and software optimization techniques. Developing and using high-quality basemaps is another way to increase performance.

6. Effect of Map Errors

Terrain-relative navigation relies on an onboard map built from orbital imagery; some error must be assumed in the building of the onboard map which causes differences in the actual locations of ground features from their predicted locations. To study the effect of map errors on our ability to meet our delivery requirement of 100m from the target, we ran example cases in our simulation in which map errors were assumed to be modeled as affine transformations (i.e., rotation about the target, translation and scaling). In the absence of map errors, the position knowledge error improves quickly to better than 100m but does not improve further after 3å images. Velocity error exhibits similar behavior. A minimum of 20 landmarks per image is desired for a good estimation of the knowledge error.

Results with simulated map errors showed that in general, the larger the map error, the greater the number of images are needed to satisfy the 100 m requirement. For instance, our simulations show that if the onboard map has a 0.5 deg rotation error combined with a 0.99 scaling error, 15 images are required to bring the knowledge error below 100 m. However, simulations with images and radar altimetry only (no velocimetry) showed that the velocity knowledge error is still in the vicinity of 4.5 m/s, which implies the need for velocimetry if there are stringent touchdown velocity requirements (e.g., < 1 m/s). Therefore, it is desirable to keep taking more images to try to keep the knowledge error as low as possible.

More research is needed to determine the usefulness of the affine transformation used to model the map errors and the degree of distortion between the onboard/"truth" maps to better assess the number of images required for Pinpoint Landing missions.

7. Radio beacon & Wind Sensing at the Landing Site

Placement of a radio beacon at the landing site is an intuitively appealing idea that is unsuitable for a first-time landing at any site due to the cost and complexity of placing the beacon. This would require a separate pathfinder mission, with its own ability to land close to a desired target and avoid hazards. However, for future missions requiring landing near a prepositioned surface asset, a radio beacon could be a useful addition. One example is a two-lander MSR mission in which one lander carries a rover which caches a sample, with a second lander landing near the sample in order to retrieve it and return it to Earth. A radio beacon could be collocated with the sample container to aid onboard navigation of the second lander. Future crewed missions obligated to land near prepositioned supplies provide another example in which a radio beacon could be useful.

Perhaps even a more useful infrastructure element at the landing site would be a wind sensor. Wind drift could be significantly reduced (resulting in fuel savings potentially in the ~200kg range for an MSL-class vehicle) if the wind direction and speed were known beforehand, allowing the lander to be retargeted to deploy its chute at a point "upwind" of the target so that wind drifts the chute much closer to the target before engine ignition. This requires a sensor capable of measuring winds in-situ from altitudes of ~10km to the surface. A wind sensor and radio beacon and wind sensor could be collocated (both on the sample container in the MSR example discussed above).

8. CONCLUSIONS

Both Pinpoint and Multipoint landing can be accomplished with a system using passive imaging. Technology developments in terrain-relative navigation, powered descent guidance, and onboard navigation are required. The principal difference between Pinpoint and Multipoint modes is the amount of fuel carried. For Pinpoint Landing, the fuel onboard must be sufficient to divert to the center of the ellipse from any point in the ellipse at ignition; the required fuel is less for Multipoint depending on the divert capability chosen by the project.

8. ACKNOWLEDGEMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

[1] A. A. Wolf, J. Tooley, S. Ploen, M. Ivanov, B. Acikmese, K. Gromov, "Performance Trades for Mars Pinpoint Landing," IEEE Aerospace Conference Proceedings, March 2006, IEEE-1661, March 2006

[2] A. Wolf, E. Sklyanskly, J. Tooley, B. Rush, "Mars Pinpoint Landing Systems Trades," AAS/AIAA Astrodynamics Specialist Conference Proceedings, AAS 07-310, August 19-23, 2007

[3] A. A. Wolf, C. Graves, R. Powell, and W. Johnson, "Systems for Pinpoint Landing at Mars," AAS/AIAA Space Flight Mechanics Conference Proceedings, AAS 04-272, February 8 – 12, 2004

[4] A. Johnson and J. Montgomery, "An Overview of Terrain Relative Navigation for Precise Lunar Landing," 2008 IEEE Aerospace Conference Proceedings, March, 2008

[5] R. Gaskell, "Landmark Navigation and Target Characterization in a Simulated Itokawa Encounter," 2005 AAS/AIAA Astrodynamics Specialist Conference Proceedings, August, 2005

[6] R. Gaskell, "Automated Landmark Identification for Spacecraft Navigation", AIAA/AAS Astrodynamics Conference Proceedings, August, 2001

[7] C. Epp and T. Smith, "Autonomous Precision Landing and Hazard Detection and Avoidance Technology (ALHAT)," 2007 IEEE Aerospace Conference Proceedings, March, 2007

[8] B. Acikmese and S.R. Ploen, "Convex Programming Approach to Powered Descent Guidance for Mars Landing," AIAA Journal of Guidance, Control, and Dynamics, Vol. 30, No. 5, 1353-1366, Sept-Oct, 2007,

[9] A. Mourikis, N. Trawny, S. Roumeliotis, A. Johnson, A. Ansar, and L. Matthies, "Vision-Aided Inertial Navigation for Spacecraft Entry, Descent, and Landing," IEEE Transactions on Robotics, Vol. 25, No. 2, 264-280, 2009

[10] R.S. Park, S. Bhaskaran, J.J Bordi, Y. Cheng, A.E. Johnson, G.L. Kruizinga, M.E. Lisano, W.M. Owen, and A.A. Wolf, "Trajectory Reconstruction of the ST-9 Sounding Rocket Experiment Using IMU and Landmark Data," 2009 AAS/AIAA Astrodynamics Specialist Conference Proceedings, AAS 09-408, August 9-13, 2009

[11] S. Boyd and L. Vandenberghe, Convex Optimization, Cambridge University Press, 2004.

[12] L. Blackmore, B. Acikmese, and D. P. Scharf, "Minimum landing error powered descent guidance for mars landing using convex optimization," AIAA Journal of Guidance, Control and Dynamics, vol. 33, no. 4, 2010.

[13] J. Carson, B. Acikmese, L. Blackmore and A. Wolf, "Capabilities of Onboard, Convex Powered-Descent Guidance Algorithms for Pinpoint and Precision Landing," IEEE Aerospace Conference, paper 1580, 2011.

[14] B. Steinfeldt, M. Grant, D. Matz and R. Braun, "Guidance, navigation, and control technology system trades for mars pinpoint landing," in Proceedings of the AIAA Guidance, Navigation and Control Conference 2008, 2008.

[14] M. P. Golombek, J. R. Matijevic, E. N. DiMaggio, and R. D. Schroeder, "Rock Size Frequency Distributions at the Mars Exploration Rover Landing Sites: Impact Hazard and Accessibility," Lunar and Planetary Science XXXIV (2003)

[15] M. P. Golombek, R. E. Arvidson, J. F. Bell III, P. R. Christensen, J. A. Crisp, B. L. Ehlmann, R. L. Fergason, J. A. Grant, A. F. C. Haldemann, T. J. Parker, S. W. Squyres, and the Athena Science Team,

[16] M. P. Golombek, B. M. Jakosky, and M. T. Mellon, "Thermal Inertia of Rocks and Rock Populations and Implications for Landing Hazards on Mars," First Landing Site Workshop for MER 2003

[17] M.P. Golombek, personal communication, 5/8/08

[18] A. Johnson, "ALHAT Field Test 3 Report, Vol 3: LIDAR Terrain Relative Navigation Algorithm Performance For Autonomous Landing & Hazard Avoidance Technology," ALHAT-9-007 Vol 3, May2010

BIOGRAPHY

Aron Wolf *is the Supervisor of the EDL Guidance, Navigation, and Control Group in the Guidance, Navigation, and Control Section at the Jet Propulsion Laboratory, where he has worked since 1982. He has led pinpoint landing study efforts at JPL since 2003. Previously, he served as mission design and navigation lead for the Mars Science Laboratory mission (2002-2005), and managed the Advanced EDL work area in the Mars Technology Program. He is also currently the Navigation Team Chief for the Stardust-NExT mission to comet Tempel-1. His previous work includes the design of orbital tour trajectories for the Galileo mission to Jupiter and the Cassini mission to Saturn. He received his undergraduate degree from Brown University and his Masters' in Aerospace Engineering from the University of Michigan.*

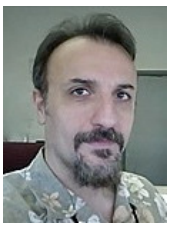

Dr. Behcet Acikmese *Behcet Acikmese received his Ph.D. in Aerospace Engineering in 2002 from Purdue University. He was a Visiting Assistant Professor of Aerospace Engineering at Purdue University before joining JPL in 2003. He is currently a technologist and a senior member of the Guidance and Control Analysis Group at Jet Propulsion Laboratory (JPL), where he is developing guidance, control, and estimation algorithms for planetary landing, formation flying spacecraft, and asteroid and comet sample return. His research interests include robust and nonlinear control, optimal control, model predictive control, convex optimization and linear matrix inequalities (LMIs) in guidance, control and estimation, and real-time optimization.*

Dr. Yang Cheng *is a senior member of the Aerial Perception Group at Jet Propulsion Laboratory (JPL), where he is developing vision algorithms for planetary spacecraft safe and precision landing. He is a key developer of numerous computer vision modules for spacecraft autonomous navigation and landing, such as the crater landmark detection/matching algorithm for spacecraft localization, the Descent Image Motion Estimation System (DIMES) for Mars Exploration Rover (MER) landing, MER visual Odometry, Rock detection and Mapping for Phoenix lander etc. He is co-PI for computer vision algorithms for safe and precise landing for a Phase A & B study of Terrain Relative Guidance System for the NASA New Millennium ST9 program. Currently, he is working on terrain relative navigation (TRN) system for future PPL missions.*

Mark C. Ivanov *is senior engineer in the EDL Guidacne, Navigation, and Control group in the Guidance, Navigation, and Control Section at the Jet Propulsion Laboratory. Mark has worked in the launch vehicle industry for over 14 years as a mission design engineer serving the Titan IV program and in numerous study efforts. Currently, Mark is involved in vehicle breakup analysis projects as well as characterizing the trade space for Mars entry studies. Mark received his undergraduate degree in aerospace engineering from the University of Florida.*

Jordi Casoliva *is an aerospace engineer in the Guidance, Navigation, and Control section at the Jet Propulsion Laboratory. He earned his M.S. degree from the University of Califonia, Irvine. Prior to joining JPL, he developed artificial intelligence control software to support several European Space Agency projects (Envisat, XMM-Newton, SMART-1). Also, he was a research staff member at "The Robotics Institute" at Carnegie Mellon University conducting research on the "Human Identification at a Distance" project. At JPL, he developed atmospheric guidance algorithms to support aerocapture, aero-gravity assist and hypersonic entry maneuvers. He also conducts performance studies to land heavy payloads on the surface of Mars, is a systems engineer in the pinpoint landing study effort, and a mission design engineer in the Supersonic Inflatable Aerodynamic Decelerator flight project. His research interests include astrodynamics, flight dynamics, atmospheric guidance, optimal control and nonlinear optimization.*

John M. Carson III *received a B.S. in Aerospace Engineering (1997) and M.S. in Engineering Mechanics (1999) from the University of Texas at Austin, followed by a Ph.D. in Mechanical Engineering (2008) from the California Institute of Technology. From 1999 to 2000 he performed flight and structural dynamics testing and analysis on the Joint Strike Fighter prototype at the Lockheed Martin Skunk Works. From 2001-present he has worked in guidance, navigation, and control design and analysis at the NASA Jet Propulsion Laboratory for several missions,*

including Cassini/Huygens and Mars Exploration Rover, and in technology development programs for small-body and planetary landing systems.